Fast Lane
Consulting and Education Services Inc.

Network Security and Design

July 2016

Background

Over 80 million records were compromised in the Anthem data breach in early 2015. In 2014, Home Depot had a staggering 109 million records compromised. According to the 2015 Verizon Data Breach Investigations Report (DBIR), there were 2,122 incidents involving loss of data. In 2014, over $70 billion had been spent on information security worldwide. These numbers undermine public confidence in an organization's efforts to protect sensitive information. Securing sensitive data requires organizational cooperation, policy enforcement, management support, and most importantly training and education.

Table of Contents

Erik Montemer

White Paper

Drivers and Motivation

Selling executive management on security spending can be a formidable challenge. A strong case must be made to justify such spending; however, careful consideration of the risks and liabilities can cause budget-conscious executives to take action. Federal, state and local laws may also drive security spending. Organizations may incur substantial fines and penalties, and possible criminal liability involving the unauthorized disclosure of sensitive information. Mastercard Worldwide established mandates requiring merchants, banks, and processors to adhere to the Payment Card Industry Data Security Standards (PCI-DSS); and may impose fines for non-compliance. Recently, C-level executives have taken notice of the financial and reputational impact of critical security incidents. Target Inc. spent "$148 million in a single quarter to cover legal fees, forensics, and other expenses"[1] incurred from a late 2013 breach involving customer data. Approximately $38 million was paid out to Target from insurance. While many organizations purchase cybersecurity and general liability insurance, the payouts may not cover the totality of spending after a major breach. The financial impact is significant to those whose information was stolen, and the organization that suffered the breach. Not every threat can be mitigated; however, risk can be reduced through proper IT infrastructure design, security policies and end-user education.

Identifying Threats

It may not be possible to predict attacks or threats, but it is possible to observe attack trends. Open Source Intelligence or OSINT is a rich and plentiful resource for tracking threats, attacks and malicious origins. Blogs, Twitter feeds, Internet Relay Chat (IRC) channels, social networks and forums can be a trove of useful intelligence. Arbor networks runs ATLAS- a "globally scoped threat analysis and monitoring system with more than 330+ ISP customers participating."[2] ATLAS collects data from global internet service providers and processes the information to generate reports on attacks, threats, and origins via geolocation. ATLAS gathers data at Layer 3 and Layer 4 of the Open Systems Interconnect (OSI) model. This includes information pertaining to specific network ports, subnets, and subnet locations. The Internet Storm Center at the SANS Institute collects similar information as ATLAS, but also tracks malware, vulnerability, and additional attack vectors. The DBIR is an exhaustive report on data breaches, security incidents, and other cybersecurity data. The DBIR reports on the prior year's data breach and incident statistics. It provides incident trends, breaches, industries affected, and most "popular" attack vectors. The report provides useful information to assist an organization in assessing their security posture.

According to the 2015 DBIR, the top three industries affected by security incidents were the public sector, financial services and information services.[3] The massive amounts of data gleaned from OSINT sources must be assimilated, evaluated, and acted upon. An organization's security posture is dependent upon carefully analyzed information. Determinations must be made as to where to focus IT security resources. Properly evaluated intelligence data may make those determinations easier. The 2015 DBIR showed that "23% of recipients now open phishing messages and 11% click on attachments."[4] This statistic indicates more end-user security training and education is needed.

Real-time threat analysis data is compiled from IT infrastructure such as firewalls, IDS/IPS appliances, server logs, endpoint security reports and various other system logs. Security Incident and Event Managers such as HP ArcSight and LogRhythm are fed system log data for the express purpose of information assimilation and evaluation. Analysts use this real-time data to undertake pro-active countermeasures; however, reliance on one data source can be problematic. It can't be overstated that information from all sources must be considered, and counter-measures developed from said data. For example, a unified threat management (UTM) appliance is reporting an increase in probes on TCP port 3389 (Remote Desktop), a review of firewall rules prohibiting TCP 3389 would be in order, in conjunction with an audit of hosts that have a business need to have TCP 3389 enabled.

Assessing Threats and Risks

Risk assessment involves analyzing "threats and vulnerabilities, impacts and likelihood"[5] of an organization's IT infrastructure. The analytical process should evaluate attack vectors - or avenues in which an attack may infiltrate the organization's network. A vulnerability analysis (also known as a vulnerability scan) examines systems for software revision levels, vulnerabilities and weak configurations. Scan results will show possible attack vectors, software revision recommendations and mitigation guidelines. It is important to determine the scope of the assessment. What systems need to be examined? What systems need not be examined?

An organization will have to determine whether to accept the risk of an unexamined system and the potential impact. Threat analysis will look at every element of risk that could conceivably happen. Threats may include natural disasters, fire, hackers, disgruntled employees and HVAC units just to name a few. Assessment data must be reviewed by the appropriate personnel: information technology, mid-level management and C-Level management (Chief Information Officer or Chief Information Security Officer). Policies and procedures must be drawn up to deal with mitigating threats and establishing controls to minimize threat risk. The aforementioned stakeholders must also determine the acceptable level of risk the organization will tolerate.

Security Policies and Framework

An information security framework must be drafted and approved by IT, human resources, legal and upper management. The framework must include the security policy, security standards, and security guidelines. The security policy must enumerate the following (this is not a comprehensive list, only some general elements):

- Acceptable use of the organization's information systems
- Prohibited use of the organization's information systems
- Password policies
- Handling of sensitive information
- User awareness of phishing and social engineering practices
- Access control to sensitive information
- Categorizing information types (i.e. confidential, sensitive, public)

The security framework must outline technical security standards. For example, Active Directory folder permissions must reflect access control to information based on department or role. The standard, in essence, dictates the baseline technical elements of the security policy. How will the organization restrict access to information? How will the organization enforce the policy's acceptable use provision?

Another example is the use of content filtering to enforce acceptable use. Filters may prohibit access to social media, explicit websites, web-based e-mail, and cloud storage such as Dropbox or Google Drive. Prohibiting access to cloud storage is a method to mitigate risk of information exfiltration. End users should already be aware that access to specific services may be filtered, because users should have been required to sign a policy acknowledgement attesting to this fact. A policy need not be excessively stringent that it prohibits users from performing their official job duties. Marketing employees may have a legitimate business need to access social media on behalf of the organization. Information security employees may have a legitimate business need to evaluate the effects of malware; in the course of this work, access to questionable websites may be required. Security standards must specify technical details such as password expiry intervals, security zones, and prohibited/permitted access between zones. The security framework is by far the most important element in network security design and operations. The infrastructure must reflect the provisions set forth in the framework. Enforcement of the framework is critical. Users must understand that any violations of the framework will have consequences. It is important that stakeholders from legal, human resources, IT, and upper management agree upon policy provisions, and it must be approved and "signed off" by C-level management.

Security Engineering

The network security infrastructure should be designed around threats and vulnerabilities, and reflect security framework requirements. Network security is the framework's primary enforcement point. When one thinks of network security most often a firewall comes to mind. Using firewalls to protect the network perimeter (or edge), today, is not sufficient. The typical "north-south" data flow (Figure 1) illustrates perimeter protection of network ingress and egress points. Even though the north-south network has evolved from its inception, it remains a relevant and practical topology for some organizations. The topology may not be overly sophisticated, or complex, but, it can be a challenge to secure.

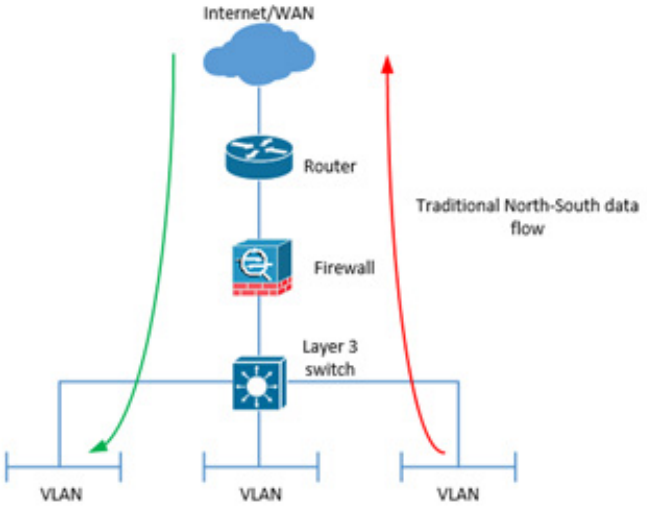

Figure 1: North-South Data Flow

Perimeter Access Control

There are a number of methods to secure the "north-south network" and the perimeter. If the router is not managed by a third party (which is typical), the organization should consider the router a first line of defense. An access control list (ACL) should be configured on the Internet ingress interface. ACLs should be general and not specific. RFC 1918 network blocks should be filtered, along with other bogus networks. One caveat to this configuration is some routers may not possess the requisite processing power to handle several hundred line ACLs. Care must be taken that the configured ACLs do not overtax the router. The subsequent line of defense is the firewall, or what has evolved into the unified threat management (UTM) appliance. UTM appliances make a perfect fit for the SMB enterprise network. A UTM appliance includes a firewall, anti-virus, anti-malware, and IDS/IPS integrated into one chassis. Generally, most UTM features are activated as subscriptions.

A firewall (and UTM) operates in what is referred to as a positive security model, also known as a “white list.” In a positive security model environment, defined traffic is permitted and everything else is rejected. Firewall policies (or rules) must follow the security standards set forth in the security framework; for instance, a defined list of restricted protocols. The list may include protocols that transmit data in clear text, such as file transfer protocols (FTP) or telnet. The list may also contain popular protocols used to exploit vulnerable systems such as Internet Relay Chat (IRC), Microsoft Windows Endpoint Mapper (TCP port 135), and Microsoft Directory Services (TCP port 445). Firewall policies must be crafted to restrict such protocols. Policies must be specific when defining source and destinations hosts. Overly broad rules—such as a /24 or larger subnet as a source or destination can increase the risk of compromised hosts. This also gives hackers and undesirables a “wide” reconnaissance source. Strict rules regarding source, destination, and ports aid in attack surface reduction. Attack surface reduction will be explained later in this paper.

Security Zone Isolation

The north-south network model may use virtual LANs (VLANs) to isolate specific parts of the network. The VLAN isolation approach is effective as an optimal network design; however, it is not always an optimal security design. It is true that each zone is isolated from the other via Layer 2 (Media Access Control), but the zones may still be reachable via Layer 3 (network).

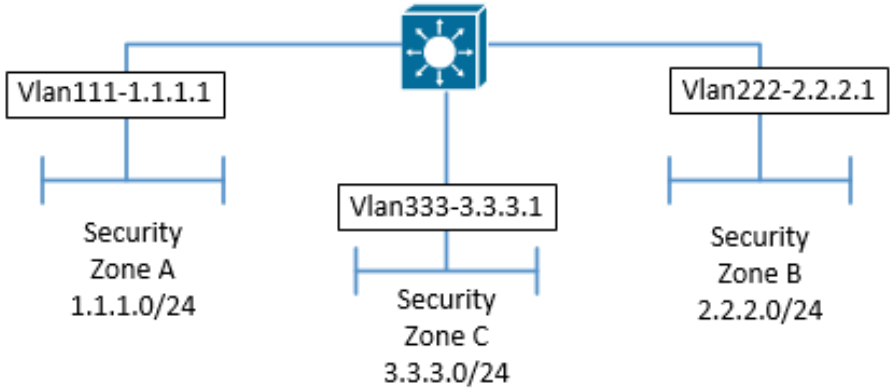

Figure 2

Each zone’s default gateway is their respective VLAN interface. These interfaces are considered connected routes in a Layer 3 switch. So, each zone can still reach the other despite them being isolated at the MAC layer. VLAN access lists can be implemented to control access between VLANs, however care must be taken not to overburden the switch or router with ACLs. A more effective, (and somewhat “audit proof”) design is illustrated in Figure 3.

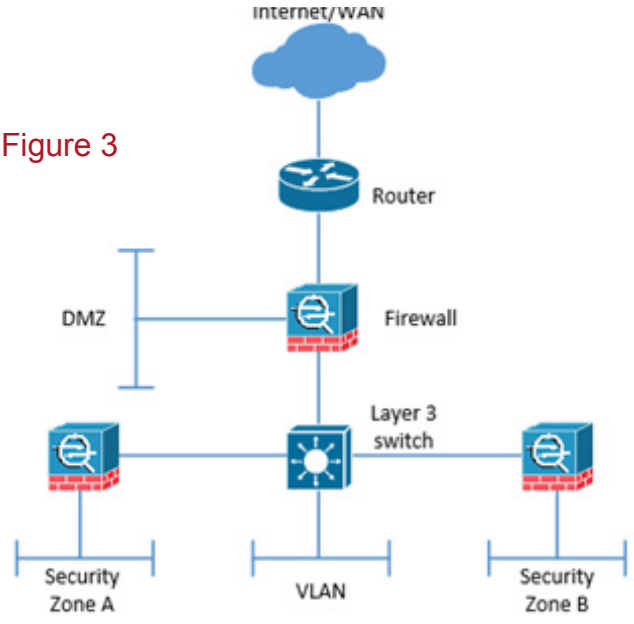

Figure 3

Security zones 'A' and 'B' are isolated from each other and other networks using firewalls. Budget conscious executives may balk at the expense of three firewalls; however, many firewall vendors have the capability to logically isolate networks within the appliance. For example, the Cisco ASA uses security contexts to isolate networks. Do not confuse security contexts with ASA interface security levels. Security levels define most secure and least secure network interfaces based on the range 0-100-0 being least secure, and 100 being most secure. Security contexts act as a "virtual ASA" within the appliance. Contexts may be defined for different zones, such as a 'Database' zone, or 'Application' zone. The network depicted in Figure 3 may use three separate firewall appliances, or three separate security contexts. It is important to refer to the security framework requirements, and security standards when designing and implementing security zones. If security zone 'A' and 'B' do not require an access control enforcement point, then VLAN isolation may suffice. Judicious use of VLAN ACLs may provide manageable access control between VLANs. A caveat to using firewalls to isolate security zones is the firewall can effectively become a router. This is not desirable as it may cause network performance problems, and place undue burden on firewall appliances.

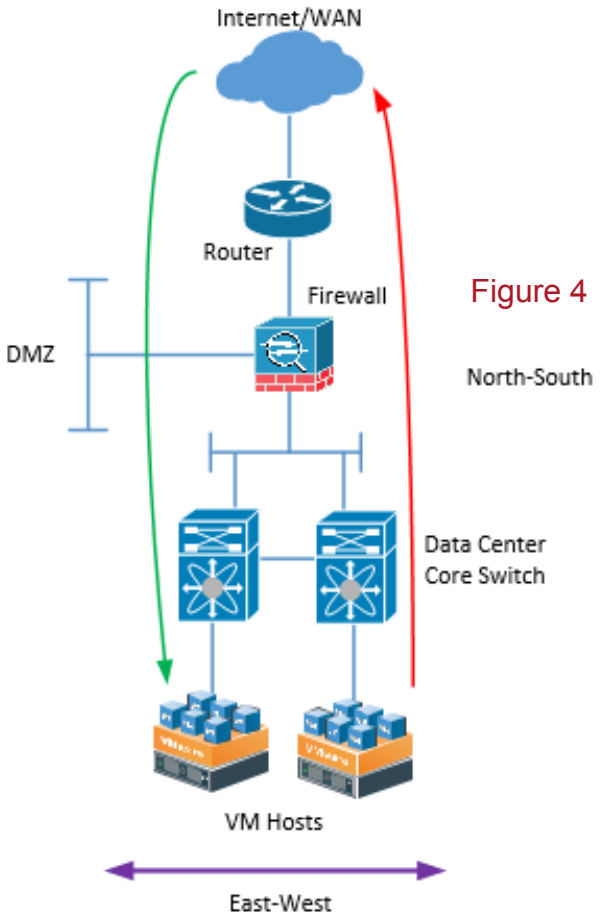

Figure 4

Applications have driven network requirements so strongly that network security has become increasingly complex. Explosive application growth has changed the face of traditional network design practices. What once was a handful of servers to support basic business functions, including databases has now morphed into the application-centric data center network. The traditional north-south network now includes east-west data flows, virtualization, and an ever-changing security posture. Figure 4 illustrates this best.

The new paradigm presents an East-West data flow, meaning server to server communications, or virtual machine to virtual machine (VM to VM) communications. This introduces a new approach to network security and an evolving security framework. Although the east-west data flow is seen more often in data center networks, it fits many topologies including branch office and smaller enterprise networks. The main driver for east-west data flow is applications. Web servers require access to database servers, and various "middleware" servers like XML gateways.

Several solutions exist to secure east-west traffic in the network. The Cisco Application Centric Infrastructure's (ACI) ASAv (Virtual ASA) acts as a Virtual Security Gateway (VSG). The ASAv combined with the Cisco Application Virtual Switch (AVS) control access at the hypervisor layer, and above. Microsegmentation is the latest use case for the ASAv and AVS. NSX is a similar solution developed by VMware. Microsegmentation controls lateral access (east-west) between virtual machines and servers. Often, an attacker will infiltrate via an ingress point in the network (the perimeter for example), whether this be a VPN termination point or ingress Internet access segment. Figure 5 illustrates the north-south network and a simple method an attacker would take to "move" around in the network laterally.

In the recent past, an attacker would use a compromised host as part of a command-and-control infrastructure. The host would connect via Internet Relay Chat (IRC) as a robot, or “bot” and a “botnet” (network of robots) would be created with varied purposes- usually to conduct Distributed Denial of Service (DDoS) attacks. Today, attacks have a more defined purpose and increased sophistication. An attacker would compromise a host, install malware, and create a malware infrastructure using compromised servers (the red server in Figure 5). The infrastructure would serve as a transport for data exfiltration. Attacks today focus on sensitive data, and data that can be sold in the black market (i.e credit card numbers, social security numbers). Microsegmentation restricts movement from a compromised host to a server and makes building a malware infrastructure much more difficult. VMWare NSX uses virtual firewalls at the hypervisor layer to restrict such movement. When deploying these solutions, an architect must continuously refer to the security standards. Specific VMs in a specific zone must either be permitted or denied on specific ports. Figure 6 illustrates this best.

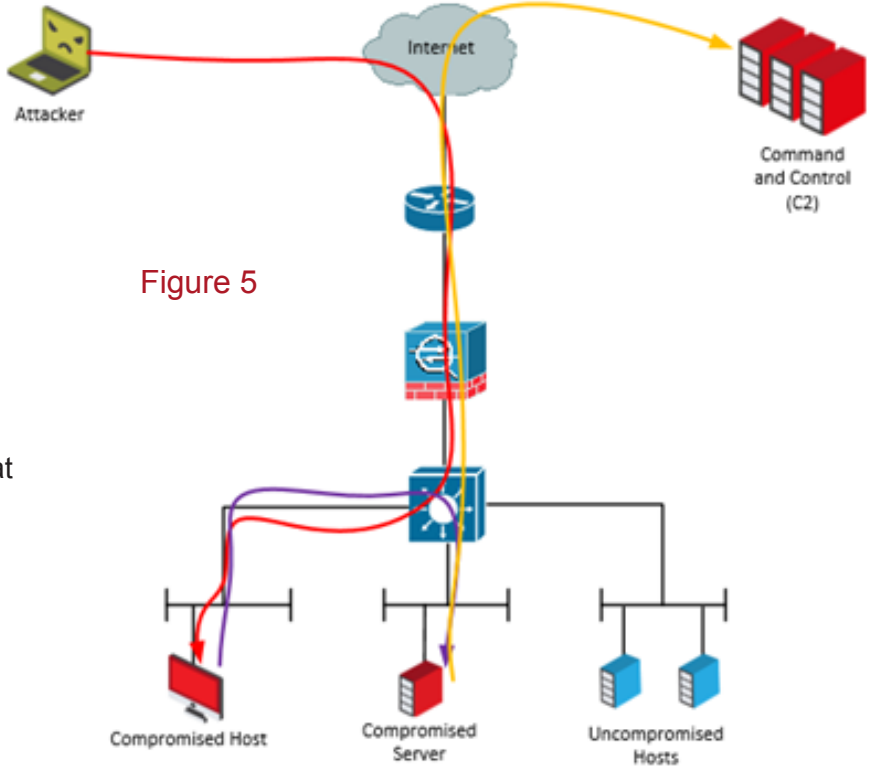

Figure 5

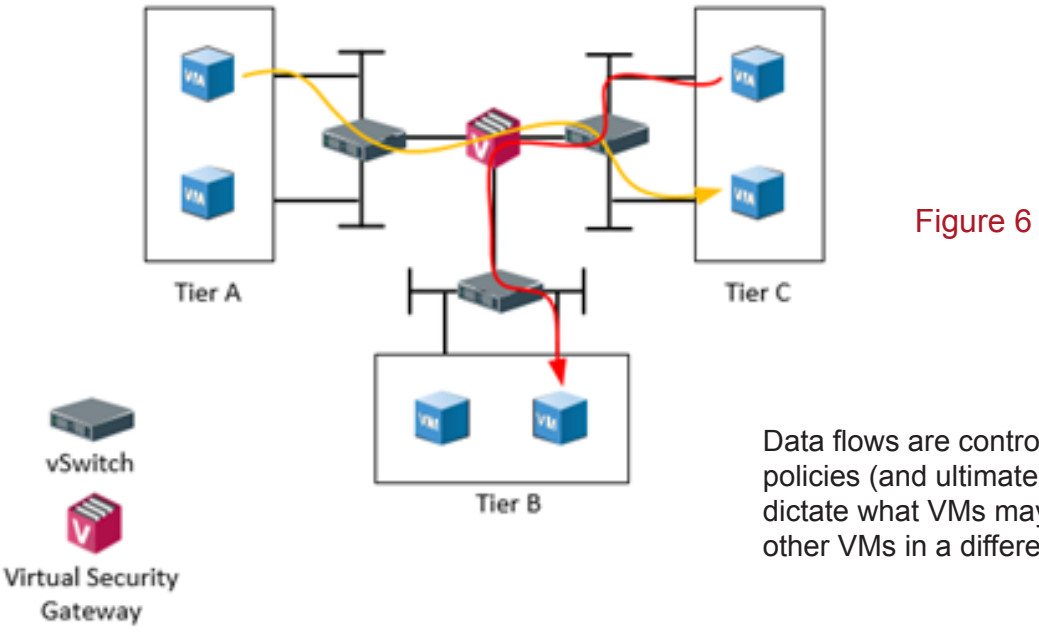

Figure 6

Data flows are controlled via the VSG. VSG policies (and ultimately the security framework) dictate what VMs may or may not connect to other VMs in a different tier.

Transit Segments and Enforcement Points

Transit segments may also be used for implementing lateral access control (east-west) or perimeter (ingress/north-south) access control. These segments are used for firewall and IDS/IPS placement in the network. These small network segments between security zones serve as an enforcement point (firewall/UTM) and a detection point (IDS/IPS and network tap) for suspicious traffic. Referring back to the north-south model network (just for argument's sake), we can see where the IDS/IPS insertion point may exist. In the model, the IDS/IPS may be placed in-line between the router and firewall. There are various implementations, and it is best to refer to the vendor guidelines as to the best practice. Some implementations rely on a switch SPAN port to monitor network traffic and alert on suspicious traffic. Another uses network taps to inspect traffic on the wire and take actions on said traffic. UTMs have an IDS/IPS capability within the appliance and thus, an optimal placement point may not be an issue. Segments may be built within a virtual switch, and a physical switch. Increased use of static routes is an important caveat to consider when building these segments.

Figure 7

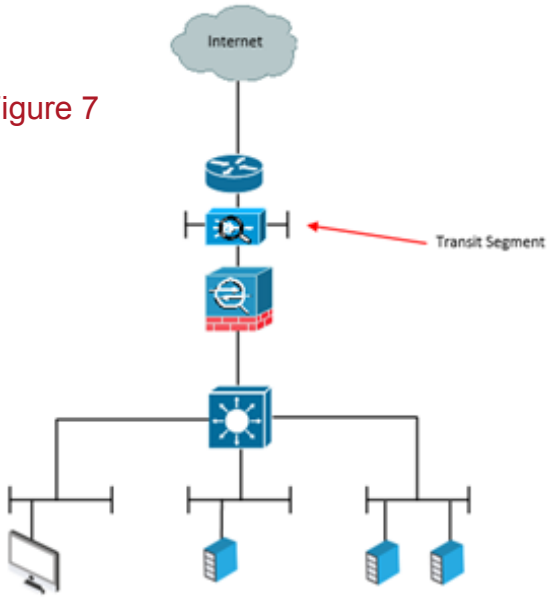

Security Auditing

It is assumed that the network is secured according to the security standards set forth in the security framework. Auditing the network involves testing the controls in the framework. For instance, the restricted protocols list includes FTP due plaintext data transmission issues. A security auditor may sit at a workstation, open a command prompt in Windows, and issue the ftp command at the prompt, and attempt to initiate an FTP transfer to a destination. If the design follows the standards and framework, the transfer should fail. The workstation should have FTP disabled as well, if not, this would be another item to add to the security standards. If one is unfamiliar with security controls, a good resource is the CIS Critical Security Controls. They can be found here: http://www.cisecurity.org/critical-controls/. Security standards may be drafted round those controls; however, controls that are specific to the business must be included as well.

Penetration testing (pen testing) is another method to test whether the security controls are adequately protecting the network. The "Threat and Risk Assessment" phase should expose any vulnerabilities and risks within the network. Remediating these vulnerabilities and threats reduces the "attack surface" of the organization. The attack surface is an organization's "reachable and exploitable vulnerabilities."[5] A thorough penetration test will determine whether attack surface reduction (reduction of "reachable and exploitable vulnerabilities) measures have been effective. The pen test will attempt to map out firewall rules and weaknesses in those rules. Additionally, the test will attempt to exploit any vulnerabilities in web servers, mail servers, or other Internet-facing hosts.

End-user workstations must also be audited. Data exfiltration (unauthorized transfer of sensitive information outside of an organization) can be prevented by establishing a security standard prohibiting active USB ports. The ports may also be configured for read-only to reduce the risk of sensitive data leaving the organization. This is only one example of several standards that should be included in the framework. Firewalls or UTMs may also be configured to prohibit access to cloud storage such as Google Drive and Dropbox.

Conclusion

Threats are continuously evolving. They've become increasingly sophisticated and require equally sophisticated methods to defend against them. Fortunately, sophisticated attacks are not always the most effective. It is thought the Anthem breach was enabled by an end-user opening a phishing e-mail. Social engineering, one of the earliest forms of hacking is still "making its rounds." It is not as effective as it once was, but is still considered a threat. Technical countermeasures cannot fix human error, but it can reduce the risk of human error. Strict e-mail filtering may reduce phishing risks. Software such as Bit9 may reduce malware risks. However, any technical solution must be backed by a comprehensive security framework. These solutions must not be haphazardly implemented.

Network security must be designed and implemented according to the framework and established security standards. Remote Desktop Protocol (RDP) remains a vulnerable protocol. How RDP and other vulnerable protocols are managed must be addressed in the security standards. The standards must dictate where the use of such protocols is acceptable. Network security devices and appliances must enforce the provisions of the framework and its associated security standards. End-users must acknowledge and agree to abide by the framework. End-user training cannot be emphasized enough. Establishing end-user security awareness programs can reduce the risk of incidents. End-users are the first line of defense, from physical security to password security, to threat awareness; the end-user may be the most effective security measure. A comprehensive end-user security program combined with a fact-based well-designed network security architecture can significantly reduce the risk of catastrophic security incidents. No one knows what the future may hold as threat actors continuously attempt to breach networks. However, cooperation, planning, and a securely engineered network may make attack attempts a waste of an attacker's time.

Sources

1 http://www.modernhealthcare.com/article/20150207/MAGAZINE/302079988
2 http://atlas.arbor.net/about/
3 Verizon DBIR pg. 3
4 Verizon DBIR pg.12
5 http://www.sans.edu/research/security-laboratory/article/did-attack-surface